

Attorney Docket No. 0023-0181

UNITED STATES PATENT APPLICATION

OF

Alireza Assadzadeh

FOR

POLICY ANALYZER

POLICY ANALYZER

TECHNICAL FIELD

[0001] The invention relates to the field of computer networks, and more particularly to systems and methods for analyzing policies in computer networks.

BACKGROUND OF THE INVENTION

[0002] Policies that control resources in a network are distributed and managed by multiple independent active system components. Currently, network design and resource policies are analyzed by setting up a test network and using various devices to inject traffic into the test network at different injection points. Designers then observe and analyze resulting traffic and behavior of the test network. If the test network does not behave as expected, designers may change the network and/or resource policies and re-inject test traffic into the network so that they may observe and analyze behavior in the changed network. Due to a lack of analysis tools, such testing is quite cumbersome, time consuming and error prone.

SUMMARY OF THE INVENTION

[0003] Systems and methods are provided for analyzing policy rules defined for a subscriber and determining packet treatment in a network.

[0004] In a first aspect of the invention, a method is provided for analyzing policy rules for a subscriber and determining packet treatment. Definitions are retrieved pertaining to policy rules for a subscriber. At least one policy point in a network is

determined based on the retrieved definitions. The packet treatment is determined at each of the at least one policy point. Information corresponding to the packet treatment is output for each of the at least one policy point.

[0005] In a second aspect of the invention, one or more devices are provided for analyzing packet treatment in a network. The one or more devices include a user input/output interface, a database interface, a management server interface, a network interface and analyzer logic. The user input/output interface is configured to receive input from a user interface and to send output to the user interface. The database interface is configured to access definitions in a database. The management server interface is configured to request and receive information from a management server. The network interface is configured to request and receive information from devices in the network. The analyzer logic is configured to analyze packet treatment based on policy rules defined for a subscriber.

[0006] In a third aspect of the invention, a system for analyzing packet treatment in a network is provided. The system includes a management server, a database, a policy analyzer and a user input/output interface. The management server is configured to load policy rules and service definitions to a router when a subscriber session is established. The database includes definitions of policy rules, service definitions and a network configuration. The database is configured to be accessible by the management server. The policy analyzer is configured to analyze packet treatment based on ones of the policy rules and the service definitions defined for a subscriber. The policy analyzer is configured to access the management server and the database. The user input/output

interface is configured to provide input to the policy analyzer and receive analysis results from the policy analyzer.

[0007] In a fourth aspect of the invention, one or more network devices are provided. The one or more network devices include an analyzer interface and a statistics module. The analyzer interface is configured to receive commands from a policy analyzer and send information to the policy analyzer. The statistics module is configured to collect statistics of traffic injected into a network. The statistics module is further configured to send the collected statistics to a policy analyzer via the analyzer interface.

[0008] In a fifth aspect of the invention, one or more devices are provided for analyzing packet treatment in a network. The one or more devices include means for receiving input from a user interface and for sending output to the user interface, means for accessing definitions in a database, means for requesting and receiving information from a management server, and means for analyzing packet treatment based on policy rules defined for a subscriber.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate an embodiment of the invention and, together with the description, explain the invention. In the drawings,

[0010] Fig. 1 illustrates an exemplary system consistent with principles of the invention;

[0011] Fig. 2 illustrates exemplary policy rules that may define a service;

- [0012] Fig. 3 is a flowchart that illustrates an exemplary process in a management server;
- [0013] Fig. 4 is a functional block diagram of an exemplary policy analyzer server;
- [0014] Fig. 5 is a functional block diagram of an exemplary policy analyzer agent;
- [0015] Fig. 6 is a functional block diagram of an exemplary management server;
- [0016] Fig. 7 is a flowchart that illustrates processing in an exemplary policy analyzer server;
- [0017] Fig. 8 is a flowchart that illustrates exemplary processing associated with initial preparations by policy analyzer server;
- [0018] Fig. 9 is a flowchart that illustrates exemplary analysis processing;
- [0019] Fig. 10 is a flowchart illustrating exemplary processing associated with determining traffic groups;
- [0020] Fig. 11 illustrates exemplary policy rules;
- [0021] Fig. 12 shows exemplary results of a traffic group analysis using the policy rules of Fig. 11;
- [0022] Fig. 13 illustrates another set of exemplary policy rules;
- [0023] Fig. 14 illustrates an entire address range and matching policy rules using the policy rules of Fig. 13;
- [0024] Fig. 15 shows traffic groups, address ranges and matching policy rules as a result of a traffic group analysis using the policy rules of Fig. 13;
- [0025] Fig. 16 illustrates the traffic groups of Fig. 15 combined into four groups of address ranges;
- [0026] Fig. 17 is a flowchart illustrating details of an exemplary “what if” analysis;

- [0027] Fig. 18 illustrates results of an exemplary “what if” analysis;
- [0028] Fig. 19 is a flowchart illustrating processing when determining traffic groups for exemplary packet/flows;
- [0029] Fig. 20 illustrates exemplary packet/flows and corresponding characteristics;
- [0030] Fig. 21 shows results of analysis using the packet/flows of Fig. 20;
- [0031] Fig. 22 shows results of traffic group analysis of the packet/flows of Fig. 20; and
- [0032] Fig. 23 is a flowchart illustrating processing in an exemplary policy analyzer server associated with injecting packets into a network and collecting packet statistics.

DETAILED DESCRIPTION

[0033] The following detailed description of the invention refers to the accompanying drawings. The same reference numbers in different drawings may identify the same or similar elements. The following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims and equivalents.

[0034] Embodiments of the invention may be implemented in hardware, software, or firmware. The firmware may be in a Read-Only Memory (ROM) and the software may reside on, for example, a medium such as a floppy disk, optical disk, or CD ROM.

Overview

[0035] Fig. 1 illustrates an exemplary system 100 consistent with principles of the invention. Exemplary system 100 may include a policy analyzer server 102, a policy analyzer user interface 104, a database 106, one or more management servers 108, policy

analyzer agents 110, subscriber devices 112, application servers 114, and a service provider network 120, including edge routers 116 and core routers 118.

[0036] Policy analyzer user interface 104 may be, for example, a browser that communicates with policy analyzer server 102 via a network. Alternatively, policy analyzer user interface 104 may be a terminal emulator directly connected to policy analyzer server 102 to communicate with policy analyzer server 102 via an asynchronous interface. Policy analyzer user interface 104 provides a user, such as a network designer, a way to interact with policy analyzer server 102.

[0037] Policy Analyzer user interface 104 may be capable of masking and layering analysis results for a display. Since the policy analysis result may be complicated and a large amount of information may be provided, policy analyzer user interface 104 may obtain input from a user about specific aspects of analysis that the user wants displayed. Each display may be modeled as a layer with a given mask. Multiple layers may be displayed together resulting in a union operation of the layers.

[0038] Policy analyzer server 102 interfaces with policy analyzer user interface 104, which may provide information to policy analyzer server 102 regarding an analysis request. Policy analyzer server 102 may send analysis results via policy analyzer user interface 104 for display to a user. Policy analyzer server 102 may interface with one or more management servers 108 and database 106 to retrieve policy rules (PRs), service definitions, subscriber definitions and a network definition. Further, policy analyzer server 102 may interface with one or more policy analyzer agents 110 to cause traffic to be injected into service provider network 120 and to measure flow of injected traffic. Policy analyzer server 102 may also interface with one or more devices in service

provider network 120 to obtain network data, which may include information regarding a state of a device and PRs installed for a subscriber.

[0039] Management servers 108 include active subscriber, service and policy definitions pertaining to service provider network 120. Management servers 108 may also provide policy, service and network data pertaining to a particular subscriber to policy analyzer server 102 when requested by policy analyzer server 102. Management servers 108 may further be responsible for downloading policy rules to appropriate routers in order to manage network traffic.

[0040] Database 106 may include one or more databases having policy rule definitions, service definitions, subscriber definitions and network definitions. Database 106 may be accessed by management servers 108 and policy analyzer server 102.

[0041] Service provider network 120 may represent a network including edge routers 116 and core routers 118. Edge routers are located at edges of network 120 near end users, such as subscribers 112 and application servers 114. Core routers 118 are located at intermediate locations within a network, such as service provider network 120. Core routers 118 may handle network traffic from several routers and typically have more high volume routing, but coarser grained policy control capability than those included in edge routers 116.

[0042] Policy analyzer agents 110 may be included in various network components, such as edge routers 116, core routers 118 and other devices, as well as in users' devices, such as in software in a subscriber's personal computer (PC), or in software in an application server. Policy analyzer agents 110 act on behalf of policy analyzer server 102 and may perform such functions as data injection, monitoring and statistics collection.

Different policy analyzer agents 110 may have different capabilities. Therefore, a subscriber's PC with special policy analyzer agent 110 software installed may be capable of performing injection, monitoring and statistics collecting services, a core router 118 may only be capable of providing monitoring and statistics collection services and an edge router 116 may only be capable of injecting certain types of packets, for example Internet Control Message Protocol (ICMP) packets.

[0043] Subscriber devices 112 refers to subscribers' network devices, such as personal computers, handheld computing devices, or any other device capable of sending and receiving information via a network.

[0044] Application server 114 provides services, such as, movies, video conferencing, voice communications, data, web pages, and the like.

Network Operation

[0045] Default policy rules and services may be defined for a subscriber based on a type of subscriber device 112 and a type of connection. For example, a subscriber accessing service provider network 120 via a dialup line, would typically not have default services defined that require broadband network access, such as video-on-demand.

[0046] A subscriber may also attempt to access a service that is not a default service. In such a case, management servers 108 may determine, based on the subscriber and service definitions stored in database 106, whether subscriber 112 is a subscriber to the service. For, example, the service definition and the subscriber definition may indicate that subscriber device 112 may be authorized to use a particular service when the subscriber agrees to pay for the service.

[0047] Fig. 2 illustrates exemplary policy rules (PRs) stored in database 106 that may define a service for a subscriber, such as a “gold” service. Gold service, in this example, is defined as PR1 and PR2. PR1 has a precedence of 30 and rate limits all File Transfer Protocol (FTP) traffic to 64 kilobits per second. PR2 has a precedence of 10 and forwards video teleconferencing traffic. In exemplary system 100, policy rules may have a precedence or priority. The lower the precedence number, the higher the priority. For example, PR2 has a higher priority than PR1 because PR2 has a lower precedence number, 10, than PR1, which has a precedence number of 30. Although this example shows higher precedence numbers as indicating lower priorities, a system may be implemented such that higher precedence numbers indicate higher priorities. In some implementations, precedence is used only when the PRs have conflicting actions.

[0048] As can be seen from this example, each PR includes a condition and an action. In PR1, the condition is FTP traffic and the action is rate limit to 64kbps. In PR2, the condition is video teleconferencing traffic and the action is forward. When traffic satisfies or matches conditions of more than one policy, one policy has a higher priority (lower precedence number) than the other matching policies and the actions cannot be combined, then the other matching policies are said to be eclipsed or overridden by the policy with the higher priority.

[0049] A particular router may define a set of eclipsing PRs for that router. The eclipsing PRs indicate incompatible PR actions. That is, if a packet satisfies conditions of multiple PRs and the corresponding PR actions are defined to be incompatible, then the router implements the highest priority PR and the corresponding action is taken. The PRs pertaining to the other actions are overridden or eclipsed. For example, a router may

define each of the following pairs of actions as incompatible (eclipsing rules):

Forward/Forward, Forward/Filter, Forward/Nexthop, Filter/Filter, Filter/Nexthop, Filter/Mark, Filter/RateLimit, Nexthop/Nexthop, Nexthop/Mark, Nexthop/RateLimit, and RateLimit/RateLimit. The actions may be defined as incompatible because performing both actions may not be possible or may not make sense. For example, it is senseless to have an action for rate limiting traffic to 64 kbps and for rate limiting the same traffic to 50 kbps simultaneously on a same router. Similarly, it may not be possible to filter traffic and to forward the same traffic on the same router.

[0050] Similarly, a particular router may have combination PR actions defined, which indicate actions that are compatible and can be combined. A combination rule indicates an order of performing actions when a packet matches more than one PR condition, regardless of the precedence or priority of the PRs. For example, a router may have the following combination PR actions defined (the second action of each pair is to be performed first): Forward/Mark, Forward/RateLimit and RateLimit/Shape with a Quality of Service (QoS). Thus, for example, suppose that a packet satisfies two PRs, PR1 and PR2. Assume PR1 has a higher priority than PR2 and an action for PR1 is Forward and an action for PR2 is RateLimit. The combination rule, in this example, defines this combination as compatible and indicates that RateLimit is to be performed first. Therefore, the action, RateLimit, will be performed first and the action, Forwarding, will be performed next, regardless of policy rule priority or precedence.

[0051] Management server 108 dynamically retrieves default policy and service definitions for subscribers from database 106 and downloads the definitions to appropriate routers 116 and/or 118 when an interface between subscriber's device 112

and router 116 is established and based on subscriber dynamic service activation requests. For example, for a Digital Subscriber Line (DSL) subscriber, this may occur when a point-to-point protocol over Ethernet (pppoe) session is started. Default policy rules and services may be determined based on policy rules, service definitions, subscriber definitions and a network definition. Further management server 108 may download additional policies to appropriate routers when a subscriber is authorized to activate and access a non-default service.

[0052] Fig. 3 is a flowchart that illustrates an exemplary process associated with management server 108 downloading default PRs and service policies to appropriate routers. One of subscriber devices 112 establishes an interface to service provider network 120 via, for example, a pppoe session. Router 116 may inform one of management servers 108 of the established interface (act 302). The management server may then retrieve a subscriber definition for a subscriber using the subscriber device (act 304). The definition may be stored in database 106 and may include type of subscriber equipment, type of connection and a list of services to which the subscriber subscribes.

[0053] The management server retrieves a network definition from database 106 (act 306). The network definition may include a configuration of service provider network 120, such as equipment types, location of equipment within source provider network 120, network addresses of equipment, versions of software loaded and executing in equipment and the like.

[0054] Next, the management server may then retrieve, from database 106, policy definitions (act 308) and, service definitions defined for the subscriber device, based on the subscriber's definition and configuration (act 310). For example, a service requiring

a high bandwidth will not be defined for the subscriber when the subscriber device is accessing the network via a low-speed dialup connection. The management server may then download the service and policy rule definitions to one or more routers 116, 118 for installation (act 312).

[0055] When the subscriber attempts to access a non-default service, the management server may then retrieve the subscriber definition, the network definition and the service definition from database 106. If the subscriber is authorized to use the service, and the service is compatible with the subscriber device, based on subscriber's definition, then the management server downloads the PRs that are defined for the service to appropriate routers 116 and/or 118 for installation on the routers.

System Components

[0056] Fig. 4 is a functional block diagram of an exemplary policy analyzer server 102. Policy analyzer server 102 includes an analyzer 402, a database interface 404, a management server interface 406, a user interface 408, an agent interface 410 and a network interface 412.

[0057] Analyzer 402 analyzes possible packet input using PRs and service definitions defined for a subscriber. The active service definitions used by analyzer 402 are based on a configuration and given service requests of subscriber devices 112.

[0058] Database interface 404 is configured to retrieve definitions from database 106. Such definitions include policy rule definitions, service definitions, subscriber definitions and a network configuration.

[0059] Management server interface 406 is configured to request and retrieve information from management servers 108, such as active policy rules, active service definitions for a subscriber, and equipment configuration of subscriber device 112.

[0060] User Interface 408 is configured to receive input, such as an analysis request, from policy analyzer user interface 104 and output information, such as analysis results, to policy analyzer user interface 104.

[0061] Agent interface 410 is configured to send commands to policy analyzer agents 110 and receive information, such as collected statistics, from policy analyzer agents 110.

[0062] Network interface 412 is configured to send and receive information to and from devices in service provider network 120, such as routers 116, 118.

[0063] Fig. 5 is a functional block diagram of exemplary policy analyzer agent 110. Policy analyzer agent 110 includes analyzer interface 502, injector 504 and statistics module 506.

[0064] Analyzer interface 502 may be configured to receive commands, such as an inject traffic command, from policy analyzer servers 102. The command may include a type of packet to inject, a packet schedule, indicating a time for injecting packets, and a packet rate. Analyzer interface 502 may further be configured to send information, such as collected statistics to one or more policy analyzer servers 102 for possible consolidation of the statistics from a group of policy analyzer agents 110.

[0065] Injector 504 may be configured to inject traffic into a service provider network 120 based on information provided from policy analyzer servers 102. The information from policy analyzer servers 102 may include a type of packet, a packet schedule, indicating a time for injecting packets, and packet rate.

[0066] Statistics module 506 may be configured to monitor injected traffic by, for example, incrementing traffic counters, measuring packet throughput, measuring network delay, counting a number of dropped packets and the like.

[0067] Although Fig. 5 shows policy analyzer agent 110 having both injector 504 and statistics module 506, some implementations of policy analyzer agent 110 may have injector 504, but not statistics module 506. Further, other implementations of policy analyzer agent 110 may have statistics module 506, but not injector 504.

[0068] Fig. 6 is a functional block diagram of management server 108. Management server 108 may include a database interface 602, a network interface 604 and a rules and service engine 606.

[0069] Database interface 602 provides access to database 106 so that management server 108 may access PR definitions, service definitions, subscriber definitions, network definitions and router definitions.

[0070] Network interface 604 provides management server 108 with an interface for downloading policy rules and network definitions to network devices, such as routers 116 and 118.

[0071] Rules and service engine 606 is configured to analyze policy rule definitions, subscriber definitions and service definitions for a subscriber, determine which policy definitions and service definitions are to be active based on the subscriber definitions, and download the active definitions to appropriate routers 116 and 118.

Policy Analyzer - Analysis

[0072] Policy analyzer user interface 104 may allow a user, such as a network designer, to interface with policy analyzer 102 and enter information for policy analyzer

analysis. For example, the network designer may enter information indicating that default policies and services pertaining to a particular subscriber are to be analyzed. Further, the network designer may enter information regarding a “what if” analysis, which will be described in detail below.

[0073] Fig. 7 is a flowchart that illustrates exemplary operation of policy analyzer server 102. Policy analyzer server 102 may perform initial preparation (act 702). For example, Fig. 8 is a flowchart that illustrates exemplary actions associated with act 702. Referring to Fig. 8, policy analyzer server 102 retrieves PR definitions from database 106 (act 802). Policy analyzer server 102 may also retrieve service definitions from database 106 (act 804). The service definitions include one or more PRs. Next, policy analyzer server 102 may retrieve a subscriber definition from database 106 (act 806). The subscriber definition includes a type of equipment and technology used by subscriber device 112 and a list of services to which a subscriber subscribes. Policy analyzer server 102 also retrieves a router definition pertaining to edge router 116 with which the subscriber communicates (act 808). The router definition may include a router type and software version as well as combination rules and eclipsing rules.

[0074] Returning to Fig. 7, policy analyzer server 102 optionally receives packet input for analysis (act 704) and retrieves subscriber and network definitions for analysis (act 706). The packet input may be used for “what if” analysis, described in detail below. Policy analyzer server 102 may also determine which policy rules and service policies to apply based on service, subscriber and network definitions, defined in database 106 (act 708). Policy analyzer server 102 also generates results based on the analysis (act 710). Policy analyzer server 102 may store the results on a machine-readable medium and/or

may provide the results to policy analyzer user interface 104 for display to the user, as described in more detail below.

[0075] For example, Fig. 9 is a flowchart that illustrates an exemplary analysis process associated with determining which policy rules and services to apply, that takes place within policy analyzer server 102. Policy analyzer server 102 accesses active service policies and default policy definitions defined for a subscriber (act 902). Policy analyzer server 102 may have retrieved this information from database 106 during initial preparation (act 702 :Fig. 7). Alternatively, service definitions and policy rules currently loaded in routers 116 and 118 for the subscriber may be retrieved from management servers 108. Policy analyzer 102 may determine one or more policy points in the network based on the retrieved definitions (act 904). A policy point is a point, such as at one of edge routers 116 or one of core routers 118 where packet behavior or treatment will be analyzed. Locations of policy points may be available in a service definition. For example, a gold gaming service, which may be a set of policy rules that define the service, may have two policy points defined: 1) on a subscriber access point and 2) on a game service access point. In many cases, policy analyzer server 102 may determine only a single policy point. For each determined policy point, policy analyzer server 102 determines packet treatment based on active policy definitions, eclipsing rules and combination rules (act 906).

[0076] For example, Fig. 10 is a flowchart that illustrates the process of determining packet treatment based on active policy definition. Fig. 10 assumes that all PRs pertaining to a subscriber are to be analyzed and resulting traffic groups are to be determined. Analysis of all possible traffic and the PRs result in one or more traffic

groups. A traffic group may represent a least common denominator of traffic that pertains to a PR or a group of PRs, as described in more detail below.

[0077] Policy analyzer server 102 may determine whether there is a portion of all possible traffic that does not satisfy any PR conditions (act 1002). If so, that portion of the traffic may be assigned to a traffic group (TG). From the remaining possible traffic, policy analyzer 102 may determine which policy condition or groups of PR conditions are satisfied and not eclipsed (act 1004). Policy analyzer server 102 may assign each portion of the remaining possible traffic to a respective TG based on the PR condition or conditions satisfied by the corresponding portion of traffic (act 1006). Policy analyzer server 102 may combine all portions of traffic into appropriate traffic groups (act 1008) and policy analyzer server 102 may store analysis results for later display (act 1010).

[0078] Returning to Fig. 9, policy analyzer server 102 may store results pertaining to packet treatment at each policy point for later display (act 908). Policy analyzer server 102 may also consolidate and correlate information pertaining to packet treatment at each policy point in order to determine packet treatment from source to destination (act 910). Policy analyzer server 102 may store the analysis results, which may show end-to-end packet treatment (act 912).

[0079] The processes described above are illustrated with an example below. Suppose exemplary PRs of Fig 11 are defined for a subscriber. PR1 is defined to rate limit FTP traffic at 64 kilobits per second and has a precedence of 30. PR2 has a precedence of 20 and is defined to forward web traffic. PR3 has a precedence of 10 and is defined to forward Internet Service Provider (ISP) content server traffic. PR4 has a

precedence of 100 and is defined to filter any traffic. Further, assume one policy point at an edge router 116 near subscriber device 112.

[0080] Referring to Fig. 10, policy analyzer server 102 determines which portions of traffic, if any, do not satisfy any of the PRs (act 1002). In this example, because all of the PRs pertain to some traffic, there is no portion of traffic that does not satisfy any of the policy rules.

[0081] Next, policy analyzer sever 102 determines which sets of remaining possible traffic satisfy one or more PRs (act 1004). In this example, FTP traffic satisfies PR1, web traffic satisfies PR2, ISP content server traffic (e.g., traffic pertaining to address 10.10.10.0/24) satisfies PR3, and any traffic satisfies PR4.

[0082] Considering the priorities or precedence and assuming that the PR actions may eclipse one another according to eclipsing rules for a router, PR3 has the highest priority (lowest precedence number), followed by PR2, PR1, and then PR4. All ISP content server traffic from address 10.10.10.0/24 has the highest priority. Therefore, policy analyzer server 102 assigns all ISP content traffic (including FTP and www traffic) to TG3 (act 1006). PR2 has the next highest priority. Therefore, policy analyzer server 102 assigns all www traffic that is not from 10.10.10.0/24 to TG1 (act 1006). PR1 has the next highest priority. Therefore, policy analyzer server 102 assigns all FTP traffic that is not from 10.10.10.0/24 to TG1. The remaining traffic, (i.e., all traffic that is not FTP or www traffic and not from 10.10.10.0/24) is assigned by policy analyzer server 102 to TG4.

[0083] Although TG1 satisfies PR1, TG2, satisfies PR2, TG3 satisfies PR3 and TG4 satisfies PR4, the different traffic groups may also be assigned in a different order, such

that, for example, TG4 matches PR1, TG3 matches PR2, TG2 matches PR3 and TG1 matches PR4.

[0084] Policy analyzer server 102 may combine the four traffic groups, if necessary (act 1008). In this example, the traffic groups do not require combining because only four types of traffic are assigned to TGs, one type of traffic for each TG. Policy analyzer server 102 may store the results for later display (act 1010).

[0085] The output of the above analysis may be displayed to the user via policy analyzer user interface 104. Fig. 12 illustrates exemplary results of the above analysis providing information about the four resulting TGs and the policies PRs satisfied by each traffic group. As illustrated in Fig. 12, each traffic group, corresponding description and policy rules with matching conditions are provided to a user, such as a network designer.

[0086] Another example is provided using addresses. Fig. 13 illustrates policies for this example. Referring to Fig. 13, PR1 is satisfied when 24 bits of a 32 bit source address match 10.10.10.0. That is, this PR covers all addresses from 10.10.10.0 – 10.10.10.255. PR2 is satisfied when 24 bits of the source address match 20.20.20.0. That is, PR2 covers all source addresses from 20.20.20.0 – 20.20.20.255. PR3 is satisfied when 16 bits of the source address match 10.10.0.0. That is, PR3 covers all addresses from 10.10.0.0 – 10.10.255.255.

[0087] Referring back to Fig. 10, policy analyzer server 102 determines portions of traffic, if any, that do not satisfy any of the PR conditions (act 1002). In this example, traffic from source addresses in a range of 0.0.0.0 – 10.9.255.255, 10.11.0.0 – 20.20.19.255, and 20.20.21.0 – 255.255.255.255 satisfy none of the PR conditions. Source address range 10.10.10.0 – 10.10.10.255 satisfies the condition of PR1, source

address range 20.20.20.0 – 20.20.20.255 satisfy the condition of PR2 and source address range 10.10.0.0 – 10.10.255.255 satisfy the condition of PR3.

[0088] Fig. 14 illustrates the entire address range and PRs satisfied by each address range. As shown in Fig. 14, address range 0.0.0.0 – 10.9.255.255 does not satisfy any policy rule conditions, address range 10.10.0.0 – 10.10.9.255 satisfies the condition of PR3, address range 10.10.10.0 – 10.10.10.255 satisfies the conditions of PR 1 and PR3, address range 10.10.11.0 – 10.10.255.255 satisfies the condition of PR3, address range 10.11.0.0 - 20.20.19.255 does not satisfy any policy rule conditions, address range 20.20.20.0 – 20.20.20.255 satisfies the conditions of PR2 and address range 20.20.21.0 – 255.255.255.255 does not satisfy any policy rule conditions.

[0089] Policy analyzer server 102 determines which sets of remaining traffic satisfy conditions of PRs (Fig. 13) and assigns the sets of traffic to traffic groups (acts 1004, 1006), as shown in Fig. 15.

[0090] Policy analyzer server 102 combines the results to form traffic groups as shown in Fig. 16 (act 1008). Fig. 16 shows the resulting four traffic groups, their address ranges and the PRs satisfied by each TG. Policy analyzer server 102 may store results (act 1010).

[0091] Referring back to Fig. 9, policy analyzer server 102 may determine packet treatment based on conditions and actions (act 906). Fig. 17 is a flowchart that illustrates processing by policy analyzer server 102 when performing a “what if” analysis to determine packet treatment at a policy point. Policy analyzer server 102 obtains a list of all PRs on an interface for a particular direction (act 1702). The directions are ingress and egress. Ingress pertains to incoming data and egress pertains to outgoing data.

Policy analyzer 102 may obtain the PRs for a particular subscriber from management server 108. Policy analyzer server 102 may also obtain the PRs for a particular subscriber's traffic from routers 116, 118.

[0092] Policy analyzer server 102 sorts the PRs based on precedence (act 1704).

Assuming that a low precedence number indicates a high priority, the sorted order may be from a low to a high precedence number. In other implementations, the sorted order may be from a high precedence number to a low precedence number.

[0093] Each PR is examined in order (act 1706). For each PR being examined, policy analyzer server 102 determines whether a given packet matches the PR condition (act 1708). If the packet matches a PR condition, the PR is added to a list of PRs that are "in effect" (act 1710). Otherwise, the PR is added to a list of PRs that are "ignored" (act 1712). Policy analyzer server 102 may also determine whether any more PRs are to be examined. If so, then policy analyzer server 102 returns to act 1706 and repeats the process.

[0094] If no more PRs are to be examined, the list of PRs that are "in effect" will be examined to determine whether any of the actions are eclipsed by other actions (act 1716). All PRs that are eclipsed, if any, are then moved to an "eclipsed" list. For example, policy analyzer server 102 may compare the actions in the "in effect" list to determine whether they appear in the eclipsing rules for the particular router 116, 118. If so, the PR that appears first eclipses other PRs. Thus, because the PRs are examined in priority order, a higher priority PR may eclipse a lower priority PR.

[0095] If two PRs are conflicting and have the same precedence, the policy analyzer server 102 may flag the PRs as invalid PRs and move them to an 'invalid' list. In another

implementation, policy analyzer server 102 may issue a “warning”. In yet another implementation, policy analyzer server 102 may use a network device definition to determine the PR that is eclipsed. If the network device definition indicates to policy analyzer server 102 that the behavior is not deterministic, policy analyzer server 102 tags both PRs as ‘in effect with warning’. In this case, a configuration order of sending the PRs to the router by management server 108 for a subscriber determines the PR that is eclipsed.

[0096] Policy analyzer server 102 may combine all actions that can be combined from the list of “in effect” PRs (act 1718). For example, “Mark” and “Forward” may be combined if they appear in the router’s combination list. The PRs are combined according to a defined order of actions in the combination rules, regardless of precedence number. The results for each PR may then be stored (act 1720) and may be displayed (act 710 :Fig. 7).

[0097] The above processing associated with Fig. 17 is illustrated by the following example. Suppose the “what if” analysis concerns a subscriber, Joe, sending packet A1 and packet A2, the edge router 116 is Juniper E-Series 5.0, Joe has a Point-to-Point Protocol (PPP) default policy, active services are be512, cp1 and cp2, direction is ingress and user IP address is 10.10.10.10. The default policy may be modeled as a special type of service with a special service session. Fig. 18 shows the service sessions, PR name, precedence, source, destination action, and analysis for this example.

[0098] Policy analyzer server 102 determines all PRs governing a policy point for the subscriber in a direction. For the ingress direction, the result is To_cp1, To_cp2, To_SAE, Internet, Cp1, and Cp2.

[0099] Policy analyzer server 102 may also sort the list of PRs based on precedence number (lowest to highest). In this example, the result is Cp1, Cp2, To_cp1, To_cp2, Internet, and To_SAE.

[00100] Policy analyzer server 102 examines each PR in order. First, policy analyzer server 102 examines packet A1 with regard to the PRs. Policy analyzer server 102 determines that packet A1 matches the condition for PR Cp1. That is, packet A1 has cp1 as its destination. Policy analyzer server 102 adds PR Cp1 to the “in effect” list. Policy analyzer server 102 determines that more PRs are to be examined.

[00101] PR Cp2 is selected for examination with respect to packet A1. Policy analyzer server 102 determines that packet A1 does not match the condition of Cp2, which is a packet having a destination of cp2. Policy analyzer server 102 adds PR Cp2 to the “ignore” list. Policy analyzer server 102 determines that there are more PRs to examine.

[00102] The analysis continues with PR To_cp1, which requires a packet with a destination of cp1 for a match. Because packet A1 satisfies this condition, policy analyzer server 102 adds PR To_cp1 to the “in effect” list. Policy analyzer server 102 determines that there are more PRs to examine.

[00103] Policy analyzer 102 examines PR To_cp2. Policy analyzer server 102 determines that packet A1 does not match the condition of PR To_cp2. Therefore, policy analyzer server 102 adds PR To_cp2 to the “ignore” list. Policy analyzer server 102 determines that there are more PRs to examine.

[00104] Policy analyzer server 102 next examines PR Internet, which requires a packet not having either a destination of cp1 or cp2. Policy analyzer server 102 determines that

packet A1 does not match this condition. Therefore, policy analyzer server 102 adds PR Internet to the “ignore” list. Policy analyzer server 102 determines that there are more PRs to examine.

[00105] Policy analyzer server 102 next examines PR To_SAE, which requires a packet having a destination of SAE (Service Activation Engine or management servers 108). Policy analyzer server 102 determines that packet A1 does not match this condition. Therefore, policy analyzer server 102 adds PR To_SAE to the “ignore” list.

[00106] Policy analyzer server 102 determines that all PRs have been examined. In this example, the “in effect” list contains [Cp1 and To_cp1] and the “ignore” list includes [Cp2, To_Cp2, Internet, and To_SAE]. Policy analyzer server 102 analyzes the actions of the “in effect” PRs with respect to eclipsing rules for router ERX 5.0. For the purposes of this example, assume the following eclipsing rules: Forward/Forward, Forward/Filter, Forward/NextHop, Filter/Filter, Filter/NextHop, Filter/Mark, Filter/RateLimit, NextHop/NextHop, NextHop/Mark, NextHop/RateLimit, RateLimit/RateLimit, and RateLimit/Filter.

[00107] Examining PR Cp1, the corresponding action is RateLimitX. The corresponding action for PR To_cp1 is NextHop. The two actions are in the eclipsing rules, as can be seen above. Because the action, PR CP1 has a lower precedence number, the corresponding action, RateLimit, eclipses the action, NextHop, for PR To_cp1. Therefore, the corresponding PR To_cp1 is moved to the eclipsed list.

[00108] The list of PRs that are in effect are examined for actions that can be combined. Currently, the “in effect” list consists of [Cp1] and the eclipsed list consists of [To_Cp1]. Only one PR is in the “in effect” list. Therefore, a combination cannot take

place. Had there been more than one PR in effect that could be combined, policy analyzer server 102 would update the “in effect” list by shifting the PRs according to the combination rules. The result is stored for each PR.

[00109] Policy analyzer server 102 performs the procedure again for packet A2, a packet to destination cp2. That is, policy analyzer server 102 retrieves the list of PRs defined for the subscriber using the interface for the ingress direction from management server 108.

[00110] Policy analyzer server 102 sorts the PRs based on precedence. The sorted PRs are [Cp1, Cp2, To_cp1, To_cp2, Internet, and To_SAE]. Policy analyzer server 102 examines each PR in order. First, policy analyzer server 102 examines PR Cp1 with respect to packet A2. The corresponding condition for PR Cp1 is a destination of cp1. Policy analyzer server 102 determines that packet A2 does not satisfy this condition. Therefore, policy analyzer server 102 adds PR Cp1 to the “ignore” list.

[00111] Policy analyzer server 102 determines that more PRs exist and examines PR Cp2. Policy analyzer server 102 determines that packet A2 matches the condition of PR Cp2, which is a packet with a destination of cp2. Policy analyzer server 102 adds PR Cp2 to the “in effect” list.

[00112] Policy analyzer server 102 determines that more PRs exist and performs act 1706 to examine PR To_cp1. Policy analyzer server 102 determines that packet A2 does not satisfy the condition of PR To_cp1, which is a packet with a destination of cp1. Policy analyzer server 102 adds PR To_cp1 to the “ignore” list.

[00113] Policy analyzer server 102 determines that more PRs exist and examines PR To_cp2. Policy analyzer server 102 determines that packet A2 satisfies the condition of

PR To_cp2, which is a packet with a destination of cp2. Policy analyzer server 102 adds PR To_cp2 to the “in effect” list.

[00114] Policy analyzer server 102 determines that more PRs exist and performs act 1706 to examine PR Internet. Policy analyzer server 102 determines that packet A2 does not satisfy the condition of PR Internet, which is a packet with a destination of not cp1 and not cp2. Policy analyzer server 102 adds PR Internet to the “ignore” list.

[00115] Policy analyzer server 102 determines that more PRs exist and examines PR To_SAE. Policy analyzer server 102 determines that packet A2 does not satisfies the condition of PR To_SAE, which is a packet with a destination of SAE. Policy analyzer server 102 adds PR To_SAE to the “ignore” list.

[00116] Policy analyzer server 102 determines that no more PRs exist. At this point, the “in effect” list consists of [Cp2 and To_cp2] and the “ignore” list consists of [Cp1, To_cp1, Internet and To_SAE].

[00117] Policy analyzer server 102 examines the PRs in the “in effect” list to determine whether any actions are eclipsed. The corresponding actions of the PRs in the “in effect” list are [RateLimit and NextHop]. The two actions are in the eclipsing rules for the router, as can be seen above. Because the action, RateLimit, corresponds to PR Cp2, which has a lower precedence number than PR To_cp2, RateLimit eclipses the action, NextHop, of PR To_cp2. Therefore, the corresponding PR, To_cp2, is moved to the eclipsed list.

[00118] The list of PRs that are in effect are examined for actions that can be combined. Currently, the “in effect” list consists of [Cp2] and the eclipsed list

consists of [To_Cp2]. Because only one PR is in the “in effect” list, no combination rules apply. As discussed above, had there been more than one PR in effect that could be combined, policy analyzer server 102 would update the “in effect” list by shifting the PRs according to the combination rules. The result is stored for each PR.

[00119] The results of the analysis may be displayed at policy analyzer user interface 104. Fig. 18 illustrates an exemplary display of the analysis results. As can be seen from Fig. 18, for packet A1, PR Cp1 is in effect, PR To_cp1 is eclipsed and PRs To_cp2, To_SAE, Internet and Cp2 are ignored. For packet A2, PR Cp2 is in effect, PR To_cp2 is eclipsed and the remaining rules are ignored. A network designer may use this information to determine how particular data/packets will be treated by service provider network 120.

[00120] As discussed above, policy analyzer server 102 may assign sets of traffic to traffic groups (Fig. 9, act 906). Fig. 19 is a flowchart that explains processing when policy analyzer server 102 analyzes given packet/flows at a particular policy point to determine traffic groups. Policy analyzer server 102 determines for each given packet/flow which policy conditions match and are in effect (act 1902). Policy analyzer server 102 determines which packet/flows belong to which traffic groups based on the effective policy rule conditions that are satisfied by the particular given packet/flow (act 1904). Policy analyzer server 102 stores the results (act 1906), which may later be displayed via policy analyzer user interface 104 (act 710 :Fig. 7).

[00121] In the following example, packet/flows are defined. The subscriber and service relationship steps are not shown. The PRs for this example are the rules shown in Fig. 11.

[00122] Fig. 20 illustrates the given packet/flows, which may be defined via policy analyzer user interface 104. A1 refers to FTP traffic from address 1.1.1.1. A2 refers to web traffic from address 1.1.1.1. A3 refers to telnet traffic from address 1.1.1.1 from source port telnet (23). A4 refers to SMTP traffic from address 1.1.1.1. A5 refers to video packet traffic from ISP content server 10.10.10.2 (server subnet 10.10.10.0/24).

[00123] Policy analyzer server 102 determines, for each packet/flow which policy conditions match and are in effect (act 1902). The process described in Fig. 17 may be used to place PRs in the "in effect" list, the "ignore" list and the eclipsed list for each packet/flow. For example, the sorted PRs may be {PR3, PR2, PR1 and PR4}.

Packet/Flow A1 is FTP traffic from 1.1.1.1. Examining the PRs (Fig. 11) in sorted order, one can see that packet/flow A1 satisfies the conditions of PRs 1 and 4. The corresponding actions are RateLimit and Filter. If we assume that the eclipsing rules are the same as in the previous examples, then the two actions have an eclipsing relationship. Because PR1 has a higher priority (lower precedence number), PR1 eclipses PR4.

[00124] Performing the same analysis for packet/flow A2, PRs 2 and 4 are satisfied, but PR 4 is eclipsed by PR2 due to precedence. Analyzing packet/flow A3, only PR4 is satisfied. Similarly, the results of the analysis of packet/flow A4, show that only PR4 is in effect for packet/flow A4, SMTP traffic from address 1.1.1.1. Packet/flow A5 is a video packet from ISP content server at 10.10.10.2 (server subnet is 10.10.10.0/24). Packet/flow A5 satisfies conditions of PR3 and PR4, but only PR3 is in effect because PR4 is eclipsed by PR3.

[00125] Fig. 21 illustrates which policies of Fig. 11 are satisfied by each packet/flow. For A1, PR1 is in effect. For A2, PR2 is in effect. For A3 and A4, PR4 is in effect. For

A5, PR3 is in effect. A network designer may use this information to determine how particular types of traffic will be handled by service provider network 120.

[00126] Policy analyzer 102 then determines to which traffic group each packet/flow belongs (act 1904). Fig. 22 illustrates the TGs with respect to packet/flows A1-A5. Referring to Fig. 22, A1 belongs to TG1, A2 belongs to TG2, A3 belongs to TG4, A4 belongs to TG 4, and A5 belongs to TG3. Policy analyzer server 102 may store the results (act 1906).

[00127] As one can see from Fig. 22, all packet/flows in a traffic group have the same PR(s) in effect. For example, A1 in TG1 has PR1 in effect, A2 in TG2 has PR2 and PR4 in effect, A5 in TG3 has PR3 in effect, and A3 and A4 in TG4 have PR4 in effect.

Packet Injection and Policy Analyzer Agents

[00128] One or more of policy analyzer agents 110 may be included in edge routers 116, core routers 118 and other equipment, such as subscribers' equipment, for example, a personal computer, handheld computer, or any other processing device. Policy analyzer agents 110 may communicate with policy analyzer server 102 and may operate under the control of policy analyzer server 102. Policy analyzer server 102 may issue commands to policy analyzer agents 110 to inject various packet types at a particular rate and may issue commands to the same policy analyzer agents 110 or other policy analyzer agents 110 to measure the results of the packet injection and to report the results to policy analyzer server 102.

[00129] When a user requests policy analysis via policy analyzer user interface 104, the user may be asked whether packets should be injected, and if so, the type and flow of

the packets to be injected. Fig. 23 is a flowchart that illustrates the process of packet injection.

[00130] Policy analyzer 102 determines injection points (act 2302). Injection points are determined based on the policy points defined in policy and service definitions, and on the flow of packets. For example, when there is one policy point and the flow is bi-directional, two injection points are needed. One injection point injects traffic from the subscriber to the content server and the other injection point injects traffic from the content server to the subscriber.

[00131] One of policy analyzer agents 110 at an injection point on a sender side injects traffic (act 2304). For each router and each policy point, the policy analyzer agent may collect injection statistics, which may indicate traffic flow, packets lost, packet rate and the like (act 2306). The collected statistics are sent to policy analyzer server 102. Policy analyzer agents 110 at intermediate locations, such as core routers 118, may collect statistics and may send their statistics to policy analyzer server 102 (act 2308). A policy agent on the receiver side sends its statistics to policy analyzer server 102 (act 2310).

[00132] Policy analyzer server 102 consolidates and correlates the statistics from each policy analyzer agent 110 (act 2312). A policy analyzer analysis performed by policy analyzer server 102 may be displayed, along with the consolidated and correlated statistics resulting from the packet injection (act 2314). These statistics may include real packet behavior/treatment that may be compared with theoretical packet behavior/treatment in service provider network 120.

[00133] The user may then analyze discrepancies between the theoretical and real packet treatment and determine a probable cause for the discrepancies (act 2316). The

reason for the discrepancies may be, for example, invalid assumptions or invalid information, such as a wrong router version number. The user may also change, via policy analyzer user interface 104, the type of data to be injected in order to determine the cause for discrepancies (act 2318). Policy analyzer server 104 and policy agent(s) 110 may again perform acts 2302 through 2316 and any discrepancies in the results may again be analyzed. The user may repeat the process as many times as desired.

Conclusion

[00134] Methods and systems consistent with the principles of the invention provide systems and methods for analyzing network traffic at policy points and for injecting traffic into the network at injection points.

[00135] The foregoing description of preferred embodiments of the invention provides illustration and description, but is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the invention for example, while a series of acts has been described with regard to Figs. 3, 7-10, 17, 19 and 23, the order of the acts may differ in other implementations consistent with the present invention. Also, non-dependent acts may be performed in parallel.

[00136] No element, act, or instruction used in the description of the present application should be construed as critical or essential to the invention unless explicitly described as such. Also, as used herein, the article “a” is intended to include one or more items. Where only one item is intended, the term “one” or similar language is used. The scope of the invention is defined by the claims and their equivalents.